arXiv:1610.05189v2 [q-bio.BM] 18 Oct 2016

ProQ3D: Improved model quality assessments using Deep Learning.

Karolis Uziela [1†], David Menéndez Hurtado [1†], Björn Wallner [2] and Arne Elofsson [1*]

[1] Department of Biochemistry and Biophysics and Science for Life Laboratory, Stockholm University, 171 21 Solna, Sweden and

[2] Department of Physics, Chemistry and Biology (IFM) /Bioinformatics. Linköping University, 581 83 Linköping, Sweden.

[*]To whom correspondence should be addressed. [†]Contributed equally

October 19, 2016

Abstract

Summary: Protein quality assessment is a long-standing problem in bioinformatics. For more than a decade we have developed state-of-art predictors by carefully selecting and optimising inputs to a machine learning method. The correlation has increased from 0.60 in ProQ to 0.81 in ProQ2 and 0.85 in ProQ3 mainly by adding a large set of carefully tuned descriptions of a protein. Here, we show that a substantial improvement can be obtained using exactly the same inputs as in ProQ2 or ProQ3 but replacing the support vector machine by a deep neural network. This improves the Pearson correlation to 0.90 (0.85 using ProQ2 input features).

Availability: ProQ3D is freely available both as a webserver and a stand-alone program at http://proq3.bioinfo.se/

Contact: arne@bioinfo.se

Supplementary information: Supplementary data are available at *arXiv* online.

1 Introduction

In 2003 we developed the first real model quality estimation program ProQ (Wallner and Elofsson, 2003). In contrast to earlier methods, such as Park and Levitt (1996), ProQ is not trained to recognise the native structure but to estimate the quality of a model. ProQ uses a machine learning approach and many features describing a protein model. In ProQ the quality is calculated for the entire model but in 2005 we extended it to estimate the quality of each residue (Wallner and Elofsson, 2006). The quality of the entire model was then estimated by summing up the predicted qualities for each residue. In ProQ2 profile weights were added to improve the predictions (Ray et al., 2012) and in ProQ3 (Uziela et al., 2016) we added energy terms calculated from Rosetta (Leaver-Fay et al., 2011). The ProQ methods have since their introduction been the best single-model based quality assessors in CASP (Kryshtafovych et al., 2016).

ProQ, ProQ2 and ProQ3 use a large number of carefully tuned inputs that are calculated from each protein model. All parameters are optimised to be independent of protein size and to have a limited range. These parameters are then used to train a support vector machine using a linear kernel (ProQ used a neural network). More advanced kernels are computationally expensive and do not produce any significant improvements. This means that ProQ2 and ProQ3 basically are linear combinations of a large set of features that all independently show a weak correlation with model quality. When these features are combined a much better correlation is achieved. However, the ProQ2 and ProQ3 methods can not identify relationships where the different features provide opposite results, i.e. it can not identify more complicated, non-linear, relationships between the features.

In the last few years machine learning using so called deep neural networks has proven to be clearly superior to other machine learning methods. These networks are able to identify non-linear relationships

between input features. We find that using identical inputs as in ProQ2 and ProQ3 but replacing the support vector machine with a deep neural network a substantial improvement can be obtained for both ProQ2 and ProQ3. The improvement is of a similar magnitude as obtained by the years of optimisation that was used to optimize the input features for ProQ3, and the gap to the consensus based assessor, Pcons (Lundstrom et al., 2001) has never been this small (CC=0.90 vs 0.95).

2 Methods

As in ProQ2 and ProQ3 a large number of features are calculated describing a model and then used to predict the quality, as measured by the S-score (Ray et al., 2012), for a single residue. Training was done using all models from CASP9 and CASP10, this is substantially more than we could use when training ProQ3. Testing was done on all models from CASP11 excluding cancelled targets and targets shorter than 50 residues. The Pearson correlation for local and global quality was used to evaluate the performance.

The learning was performed using the Keras Python library with the Theano backend. We used two dense hidden layers with 200 and 600 neurons respectively. Increasing the number of layers and neurons did not improve the results. The final model was trained with Adadelta and 10^{-11} penalty for the L^2 regularization and shuffling the training data, for details see supporting information.

3 Results and Discussion

To estimate the global quality of a protein model with ProQ the predicted qualities for each residue are summed and the sum is divided by the protein sequence length. The correlation of overall (global) quality is substantially improved using either ProQ2 or ProQ3 inputs both when we calculate the correlation for all models together (CC-glob) or when the per target correlations are studied (CC-target), Table 1 and Figure 1. In addition, protein model quality assessment can be used to identify good and bad regions of a model, i.e. the local quality. Here we obtain a similar improvement as achieved for the global

Table 1: **Performance of different QA methods on CASP11.**

	CC-glob	CC-target	CC-loc	CC-model	GDT_loss
ProQ	0.60	0.44	0.50	0.39	0.06
Qprob	0.71	0.56	-	-	0.07
Qmean	0.73	0.57	0.57	0.42	0.08
ProQ2	0.81	0.65	0.69	0.47	0.06
ProQ3	0.85	0.65	0.73	0.51	0.06
ProQ2D	0.85	0.68	0.72	0.49	0.05
ProQ3D	0.90	0.71	0.77	0.54	0.06
	Consensus based methods				
Pcons	0.95	0.77	0.87	0.68	0.07

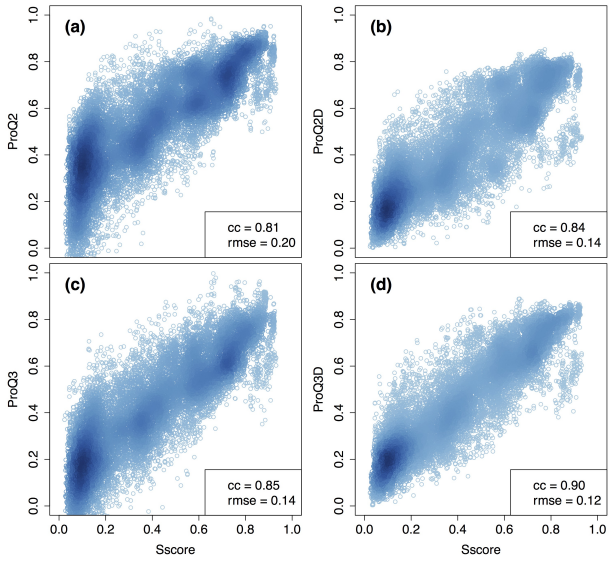

Figure 1: Predicted global model quality on CASP11 data set for ProQ2 (top) and ProQ3 (bottom) trained with a SVM (left) or Theano (right). RMSE is a root mean square error between the prediction and the target value (lower values are better). (a) ProQ2, SVM, (b) ProQ2D, Theano, (c) ProQ3, SVM, (d) ProQ3D, Theano

quality estimation independently if it is measured for all residues (CC-loc) or per model (CC-model) Table 1.

Finally, we examined if better top-ranked models could be selected using the new quality estimators. We calculated the average GDT loss of the first ranked models for each method. Unfortunately, the selection of top ranked models does not show any significant improvement between any of the top QA methods. However, the same is observed when using Pcons, although the correlation is probably quite close to the theoretical limit, as it is higher than the correlation between different methods used to evaluate the quality of a model Wallner and Elofsson (2007). This indicates that to advance further it will be necessary to use a different approach.

Funding

This work was supported by grants from the Swedish Research Council (VR-NT 2012-5046 to AE and 2012-5270 to BW) and Swedish e-Science Research Center (BW). Computational resources were provided by the Swedish National Infrastructure for Computing (SNIC) at NSC.

References

Kryshtafovych, A., Barbato, A., Monastyrskyy, B., Fidelis, K., Schwede, T., and Tramontano, A., 2016. Methods of model accuracy estimation can help selecting the best models from decoy sets: Assessment of model accuracy estimations in CASP11. *Proteins*, **84 Suppl 1**:349–369.

Leaver-Fay, A., Tyka, M., Lewis, S., Lange, O., Thompson, J., Jacak, R., Kaufman, K., Renfrew, P., Smith, C., Sheffler, W., *et al.*, 2011. ROSETTA3: an object-oriented software suite for the simulation and design of macromolecules. *Methods Enzymol*, **487**:545–574.

Lundstrom, J., Rychlewski, L., Bujnicki, J., and Elofsson, A., 2001. Pcons: a neural-network-based consensus predictor that improves fold recognition. *Protein Sci*, **10**(11):2354–2362.

Park, B. and Levitt, M., 1996. Energy functions that discriminate x-ray and near native folds from well-constructed decoys. *J Mol Biol*, **258**(2):367–392.

Ray, A., Lindahl, E., and Wallner, B., 2012. Improved model quality assessment using proq2. *BMC Bioinformatics*, **13**:224.

Uziela, K., Shu, N., Wallner, B., and Elofsson, A., 2016. Proq3: Improved model quality assessments using rosetta energy terms. *Sci Rep*, **6**:33509.

Wallner, B. and Elofsson, A., 2003. Can correct protein models be identified? *Protein Sci*, **12**(5):1073–1086.

Wallner, B. and Elofsson, A., 2006. Identification of correct regions in protein models using structural, alignment, and consensus information. *Protein Sci*, **15**(4):900–913.

Wallner, B. and Elofsson, A., 2007. Prediction of global and local model quality in CASP7 using pcons and proq. *Proteins*, **69 Suppl 8**:184–193.

arXiv:1610.05189v2 [q-bio.BM] 18 Oct 2016

Supporting material for: "ProQ3D: Improved model quality assessments using Deep Learning."

Karolis Uziela, David Menendez Hurtado, Nanjiang Shu, Björn Wallner and Arne Elofsson

1 Supplementary Results

Here, we provide results that are analogical to the ones in the main text but on another data set - CAMEO (Haas *et al.*, 2013). All CAMEO models from 2014-06-06 to 2015-05-30 were used. Table S1 shows the correlations and GDT_loss for CAMEO data set and Figure S1 shows the model quality prediction scatter plots. As we already noted for CASP11, the deep learning versions of the predictors (ProQ2D and ProQ3D) are superior to their SVM-based counterparts (ProQ2 and ProQ3).

Additionally, we provide the local prediction (residue-level) scatter plots for CASP and CAMEO data sets (Figure S2 and S3). The density plots suggest that the deep learning versions of methods are better at distinguishing low quality models on CASP11 data set (S-score < 0.1) and high quality models on CAMEO data set (S-score > 0.9).

In the benchmark we used the updated ProQ version (ProQres) that is able to predict the local model quality (Wallner and Elofsson, 2006). We have also compared our methods with two other state-of-the-art methods: Qmean (Benkert *et al.*, 2008) and Qprob (Cao and Cheng, 2016), as well as a reference consensus method Pcons (Lundstrom *et al.*, 2001).

2 Methods

2.1 Architecture

The chosen architecture is that of a Multi Layer Perceptron (MLP) (LeCun *et al.*, 1998) with two hidden layers of 600 and 200 hidden units respectively, and using ReLU (Glorot *et al.*, 2011) as a non linearity. The parameters are optimised with Adadelta (Zeiler, 2012), a modification on the Stochastic Gradient Descent that can automatically adapt the learning rate, trying to optimise the mean squared error. We used two regularisation methods: dropout (Srivastava *et al.*, 2014) and L^2 penalty on the weights. Varying the number of neurons in each layer did not yield a significant difference in the results.

Adding more hidden layers gives worse results because the convergence is slowed down (Glorot and Bengio, 2010), and the model does not gain expressiveness since it is already an universal approximator (Hornik *et al.*, 1989).

2.2 Dropout

With dropout, at training time, we set to zero ("drops") a random fraction p of the hidden neurons of each layer in each iteration, while multiplying the rest by a factor $1/p$. For a model with H hidden units, the output at test time is the geometric mean of the 2^H possible combinations, each of which has seen at most one data point, and are strongly regularised towards the consensus. Following the original paper, we set $p = 0.5$.

2.3 L^2 penalty

With the aim to discourage large weights and improve generalisation accuracy, we add a term to the loss function:

$$\lambda \sum W_{ij}^2, \tag{1}$$

where λ is a scalar and W are the weight matrices of the MLP.

This prevents weights to grow too large unless they are sufficiently beneficial, that is, they can overcome the penalty in the loss function.

We tested different training penalties ranging from 10^{-4} to 10^{-12}, with minor effects on the quality of the results.

2.4 Implementation

The implementation was done using the Keras framework (Chollet, 2015) and the Theano backend (Theano Development Team, 2016), and trained on GPUs. However, the ProQ3D program does not require GPU to run, because prediction works fast enough on CPU (less than 1 second per model, excluding feature generation).

Table S1: **Performance of different QA methods on CAMEO data set**

	CC-glob	CC-target	CC-loc	CC-model	GDT_loss
ProQ	0.67	0.54	0.52	0.45	0.04
Qprob	0.68	0.53	-	-	0.04
Qmean	0.68	0.48	0.54	0.47	0.05
ProQ2	0.75	0.57	0.60	0.50	0.04
ProQ3	0.79	0.58	0.65	0.55	0.04
ProQ2D	0.79	0.60	0.64	0.52	0.04
ProQ3D	0.82	0.60	0.69	0.58	0.04
	Consensus based methods				
Pcons	0.89	0.67	0.81	0.70	0.05

References

Benkert, P., Tosatto, S., and Schomburg, D. (2008). QMEAN: A comprehensive scoring function for model quality assessment. *Proteins*, **71**(1), 261–277.

Cao, R. and Cheng, J. (2016). Protein single-model quality assessment by feature-based probability density functions. *Sci Rep*, **6**, 23990.

Chollet, F. (2015). Keras. https://github.com/fchollet/keras.

Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS10). Society for Artificial Intelligence and Statistics*.

Glorot, X., Bordes, A., and Bengio, Y. (2011). Deep sparse rectifier neural networks. In G. J. Gordon and D. B. Dunson, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS-11)*, volume 15, pages 315–323. Journal of Machine Learning Research - Workshop and Conference Proceedings.

Haas, J., Roth, S., Arnold, K., Kiefer, F., Schmidt, T., Bordoli, L., and Schwede, T. (2013). The protein model portal–a comprehensive resource for protein structure and model information. *Database (Oxford)*, **2013**, bat031.

Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, **2**(5), 359 – 366.

LeCun, Y., Bottou, L., Orr, G. B., and Müller, K.-R. (1998). Effiicient backprop. In *Neural Networks: Tricks of the Trade, This Book is an Outgrowth of a 1996 NIPS Workshop*, pages 9–50, London, UK, UK. Springer-Verlag.

Lundstrom, J., Rychlewski, L., Bujnicki, J., and Elofsson, A. (2001). Pcons: a neural-network-based consensus predictor that improves fold recognition. *Protein Sci*, **10**(11), 2354–2362.

Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, **15**(1), 1929–1958.

Theano Development Team (2016). Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, **abs/1605.02688**.

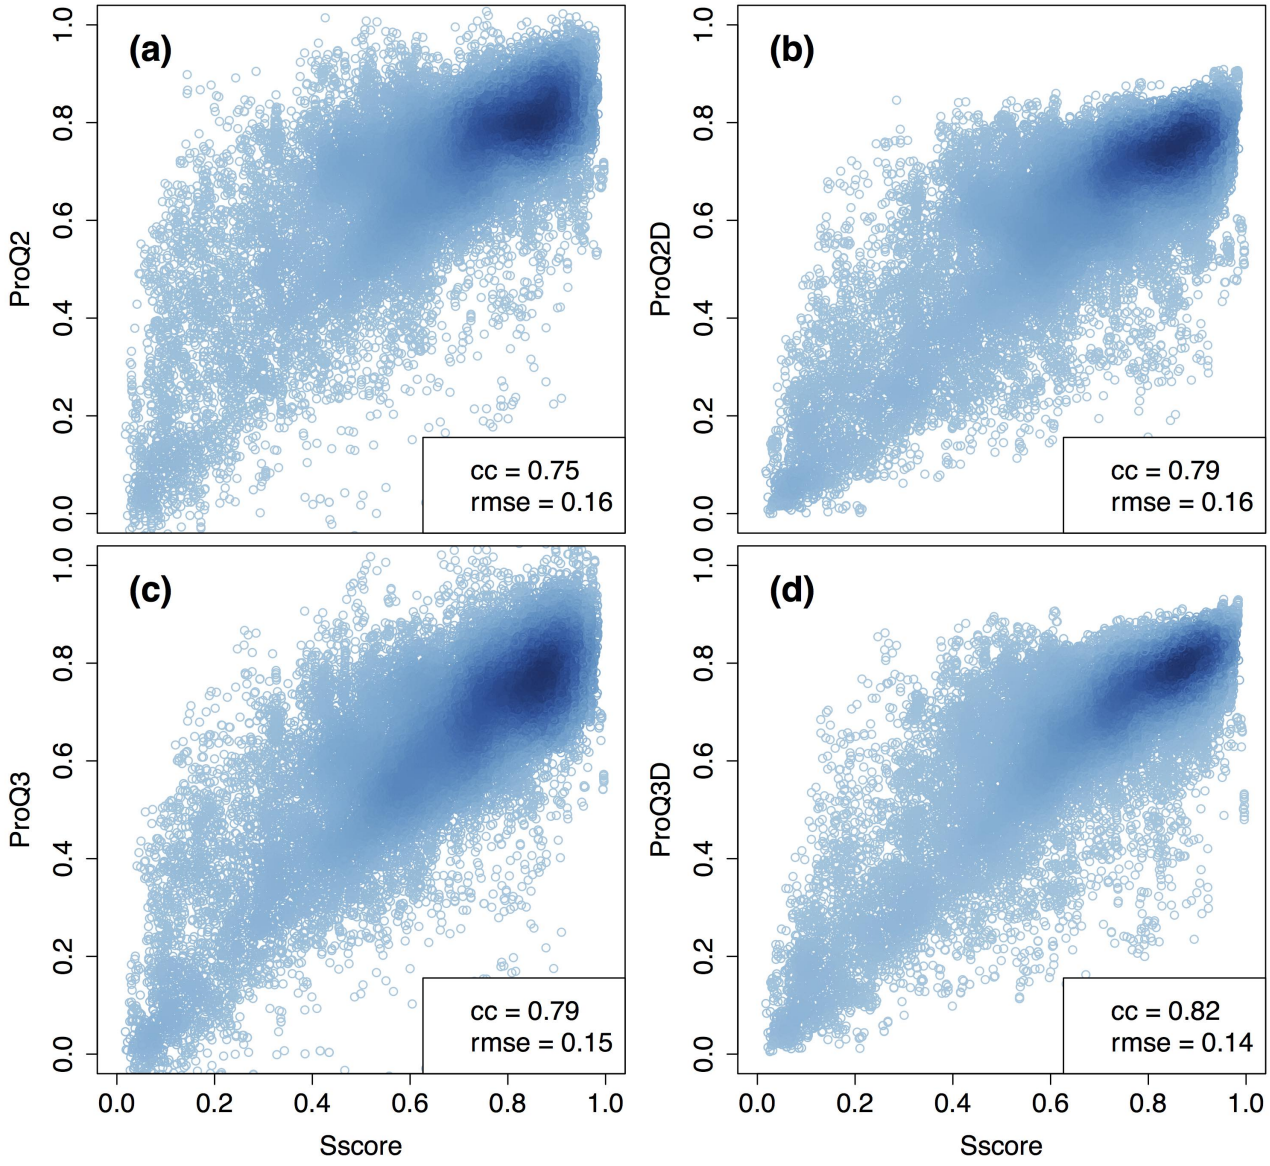

Figure S1: Predicted global model quality on CAMEO data set for ProQ2 (top) and ProQ3 (bottom) trained with a SVM (left) or Theano (right). RMSE is a root mean square error between the prediction and the target value (lower values are better). (a) ProQ2, SVM, (b) ProQ2D, Theano, (c) ProQ3, SVM, (d) ProQ3D, Theano

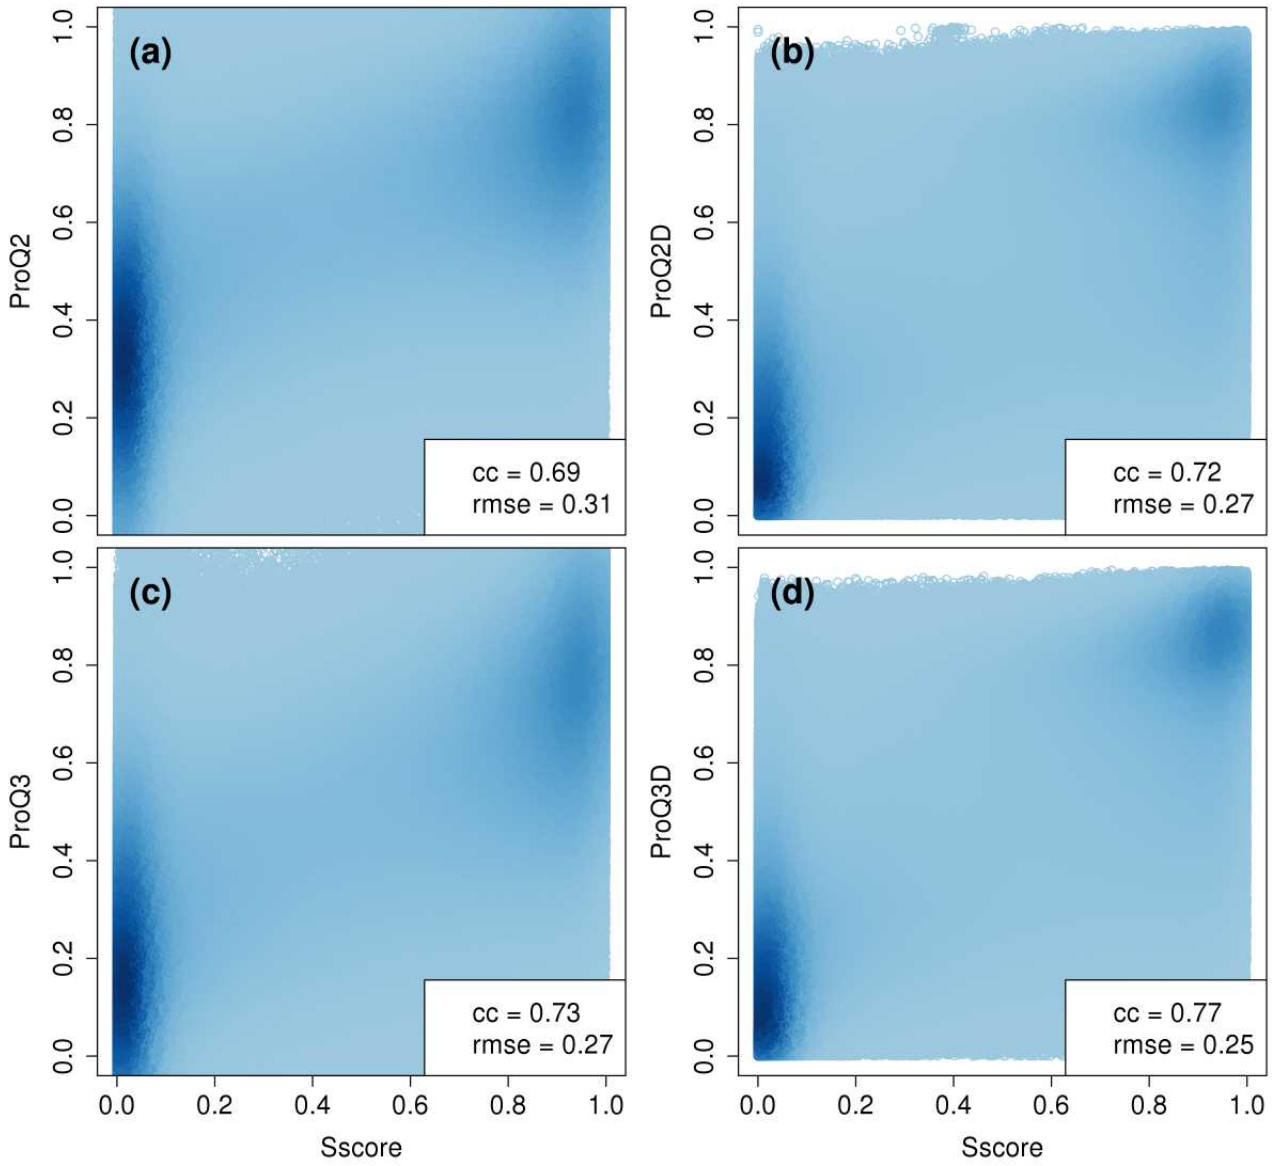

Figure S2: Predicted local model quality on CASP11 data set for ProQ2 (top) and ProQ3 (bottom) trained with a SVM (left) or Theano (right). RMSE is a root mean square error between the prediction and the target value (lower values are better). (a) ProQ2, SVM, (b) ProQ2D, Theano, (c) ProQ3, SVM, (d) ProQ3D, Theano

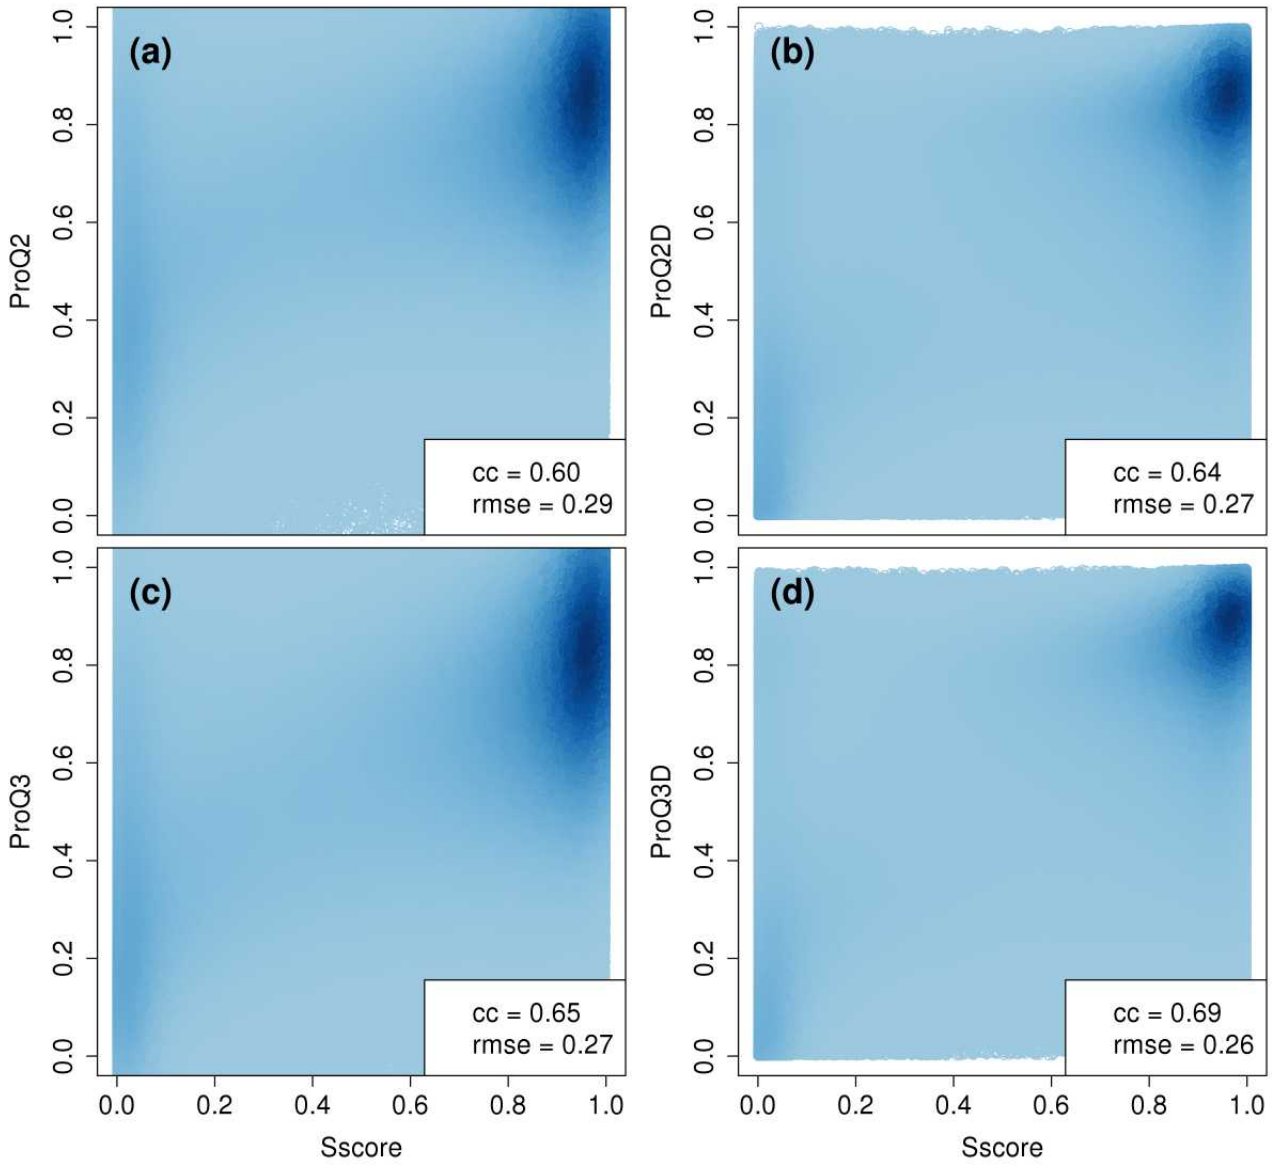

Figure S3: Predicted local model quality on CAMEO data set for ProQ2 (top) and ProQ3 (bottom) trained with a SVM (left) or Theano (right). RMSE is a root mean square error between the prediction and the target value (lower values are better). (a) ProQ2, SVM, (b) ProQ2D, Theano, (c) ProQ3, SVM, (d) ProQ3D, Theano

Wallner, B. and Elofsson, A. (2006). Identification of correct regions in protein models using structural, alignment, and consensus information. *Protein Sci*, **15**(4), 900–913.

Zeiler, M. D. (2012). ADADELTA: an adaptive learning rate method. *CoRR*, **abs/1212.5701**.